

OSG Technologies Updates

Brian Bockelman
OSG AHM 2014

This presentation

- I'll cover topics from several OSG functional areas, including:
 - Technology (and software): with inputs from Tim Cartwright and Tim Theisen.
 - Campus Grids: with inputs from Rob Gardner.
 - Security: with inputs from Mine Altunay.
- Thanks to all those who contributed slides!

OSG Technology

- OSG software allows the OSG and sites to advance the science of DHTC.
- A few thrusts of the next year:
 - Continue the stable running of the OSG Software stack.
 - Make significant gains in usability.
 - Deploy new technologies into the software stack.
 - Incorporate new use cases into the software stack.

The Software Factory

- One important service OSG provides is a “software factory”.
- Raw components (software packages) go in one side. A software distribution comes out the other.
- We assemble/integrate the components, improve them, test them, and distribute the results to the OSG.
- We are also developing a “Software Factory Factory”; other organizations, such as HTCondor, HCC, I2, and USCMS are investigating how to use our infrastructure to produce distributions.

Example: HTCondor

- OSG adds 8 patches, e.g.:
 - Patch start/stop script to get OSG security values
 - Ensure proxies are ≥ 1024 bits (contributed back)
- Integrated with other packages, e.g.:
 - Globus GRAM gatekeeper as batch system
 - GlideinWMS pilot jobs and central manager
- Automated tests include:
 - “Regular” HTCondor job
 - HTCondor-G job \rightarrow GRAM \rightarrow HTCondor backend
- We contributed unified source RPM to CHTC

Software Releases

	Q1	Q2	Q3	Q4
Year 1	5	3	3	4
Year 2	4	4 / 1	4 / 5	—

- Now on a predictable monthly schedule
- Extra releases for security or critical updates
 - Jun 2013: CA certificates (5 days)
 - Dec 2013: React to OS changes (9 days)
 - Feb 2014: Critical OSG 3.2 update (3 days)
- Tickets closed: 423 last year, 365+ this year

Software Release Series

- OSG has a new release series, 3.2.x. Release series boundaries give us a chance to remove obsolete components and package disruptive upgrades (HDFS).
 - OSG 3.0 -> OSG 3.1: 44% increase in number of RPMs.
 - OSG 3.1 -> OSG 3.2: 15% decrease in number of RPMs. About 25% of RPMs are identical to those in EPEL (and have a minimal support load).
- I believe this release series will run for >2 years. We will add support for RHEL7 without doing a new series (unlike 3.0 to 3.1). When we do release 3.3, I hope to have another 20% decrease in the number of RPMs.
- Any newly requested packages (such as xrootd4) will go into 3.2 only.

Software Maintenance

- The Software Team had some major maintenance challenges it tackled in the last year:
 - **SHA-2 support:** New suite of encryption algorithms. Required a complete revalidation of all security-related components. Required OSG to write significant patches to JGlobus and BestMan.
 - **Java:** Moved from Oracle JDK 6 to OpenJDK 7. Required a complete revalidation of the Java components.
 - **OpenSSL upgrade:** RHEL6.5 included a major upgrade to OpenSSL which broke several grid components.

Gains In Usability

- OSG has always had a thin middleware layer (for some value of “thin”); the user-friendly interfaces were always expected to come from VOs.
- Many data points in the past (early RSV) and recently (BOSCO) show that OSG continues to struggle with producing user-friendly products.
- Current focus is on improving services and reducing barriers, not new products.

New Service - OSG Connect

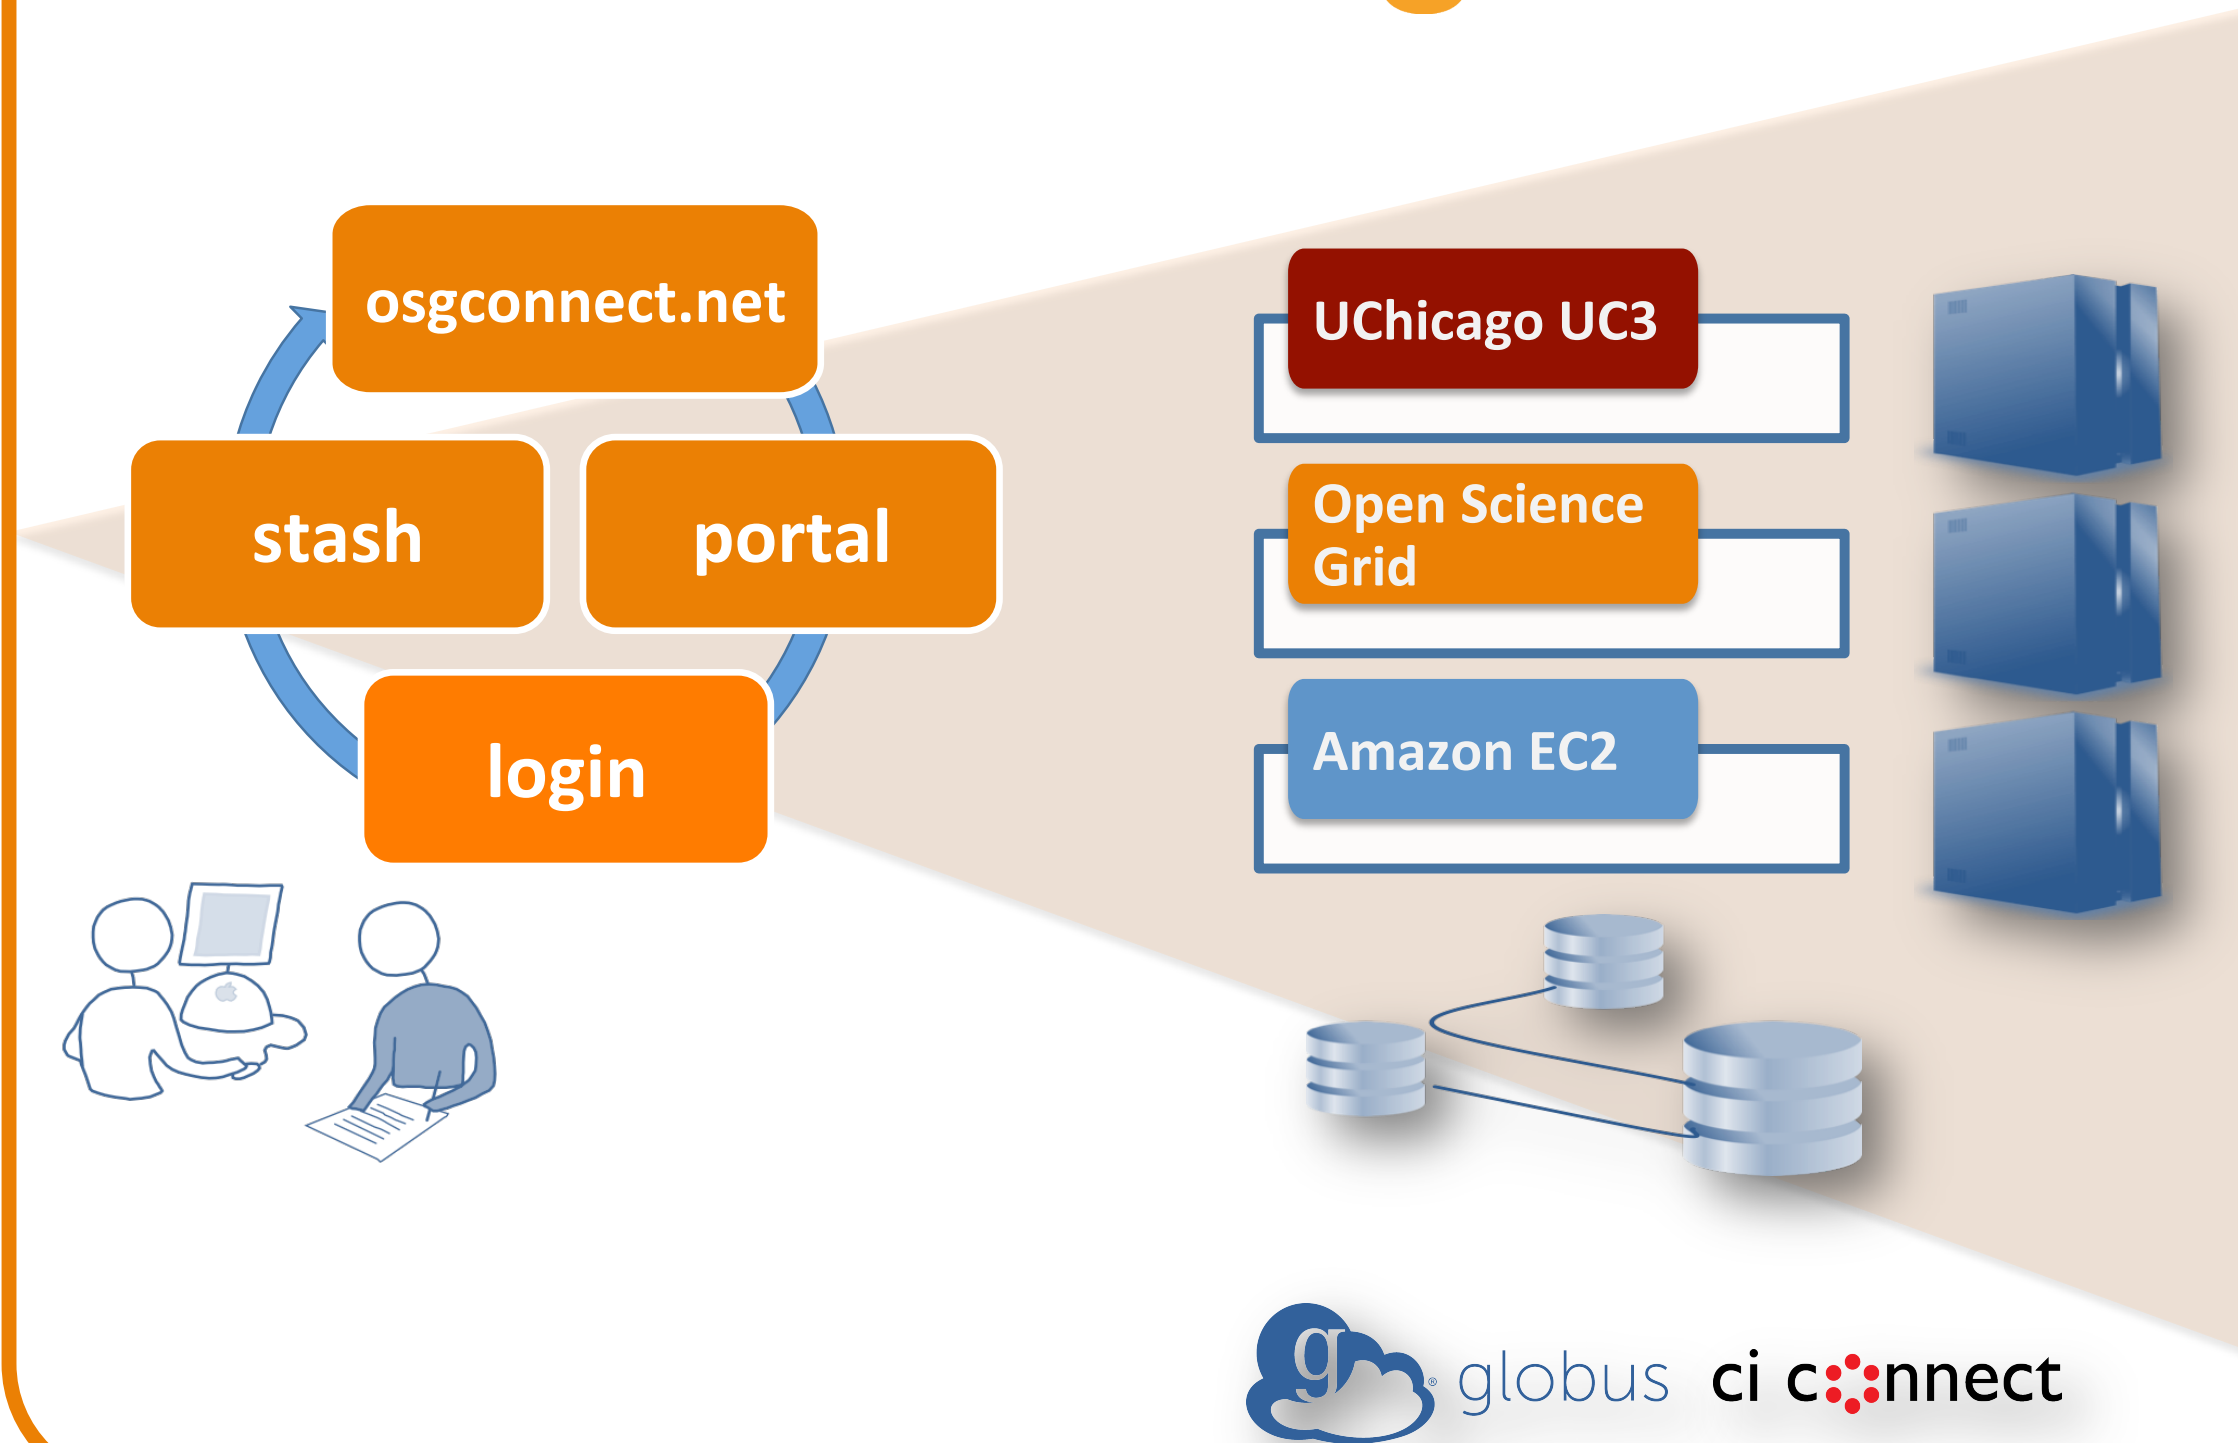
- OSG Connect, from the Campus Grids Area, is a new service to bootstrap a new DHTC user.
- <http://osgconnect.net>
- Idea is that individuals can start running jobs within 30 minutes; no software install needed.
- Further, OSG will run a instance as a service for a campus.

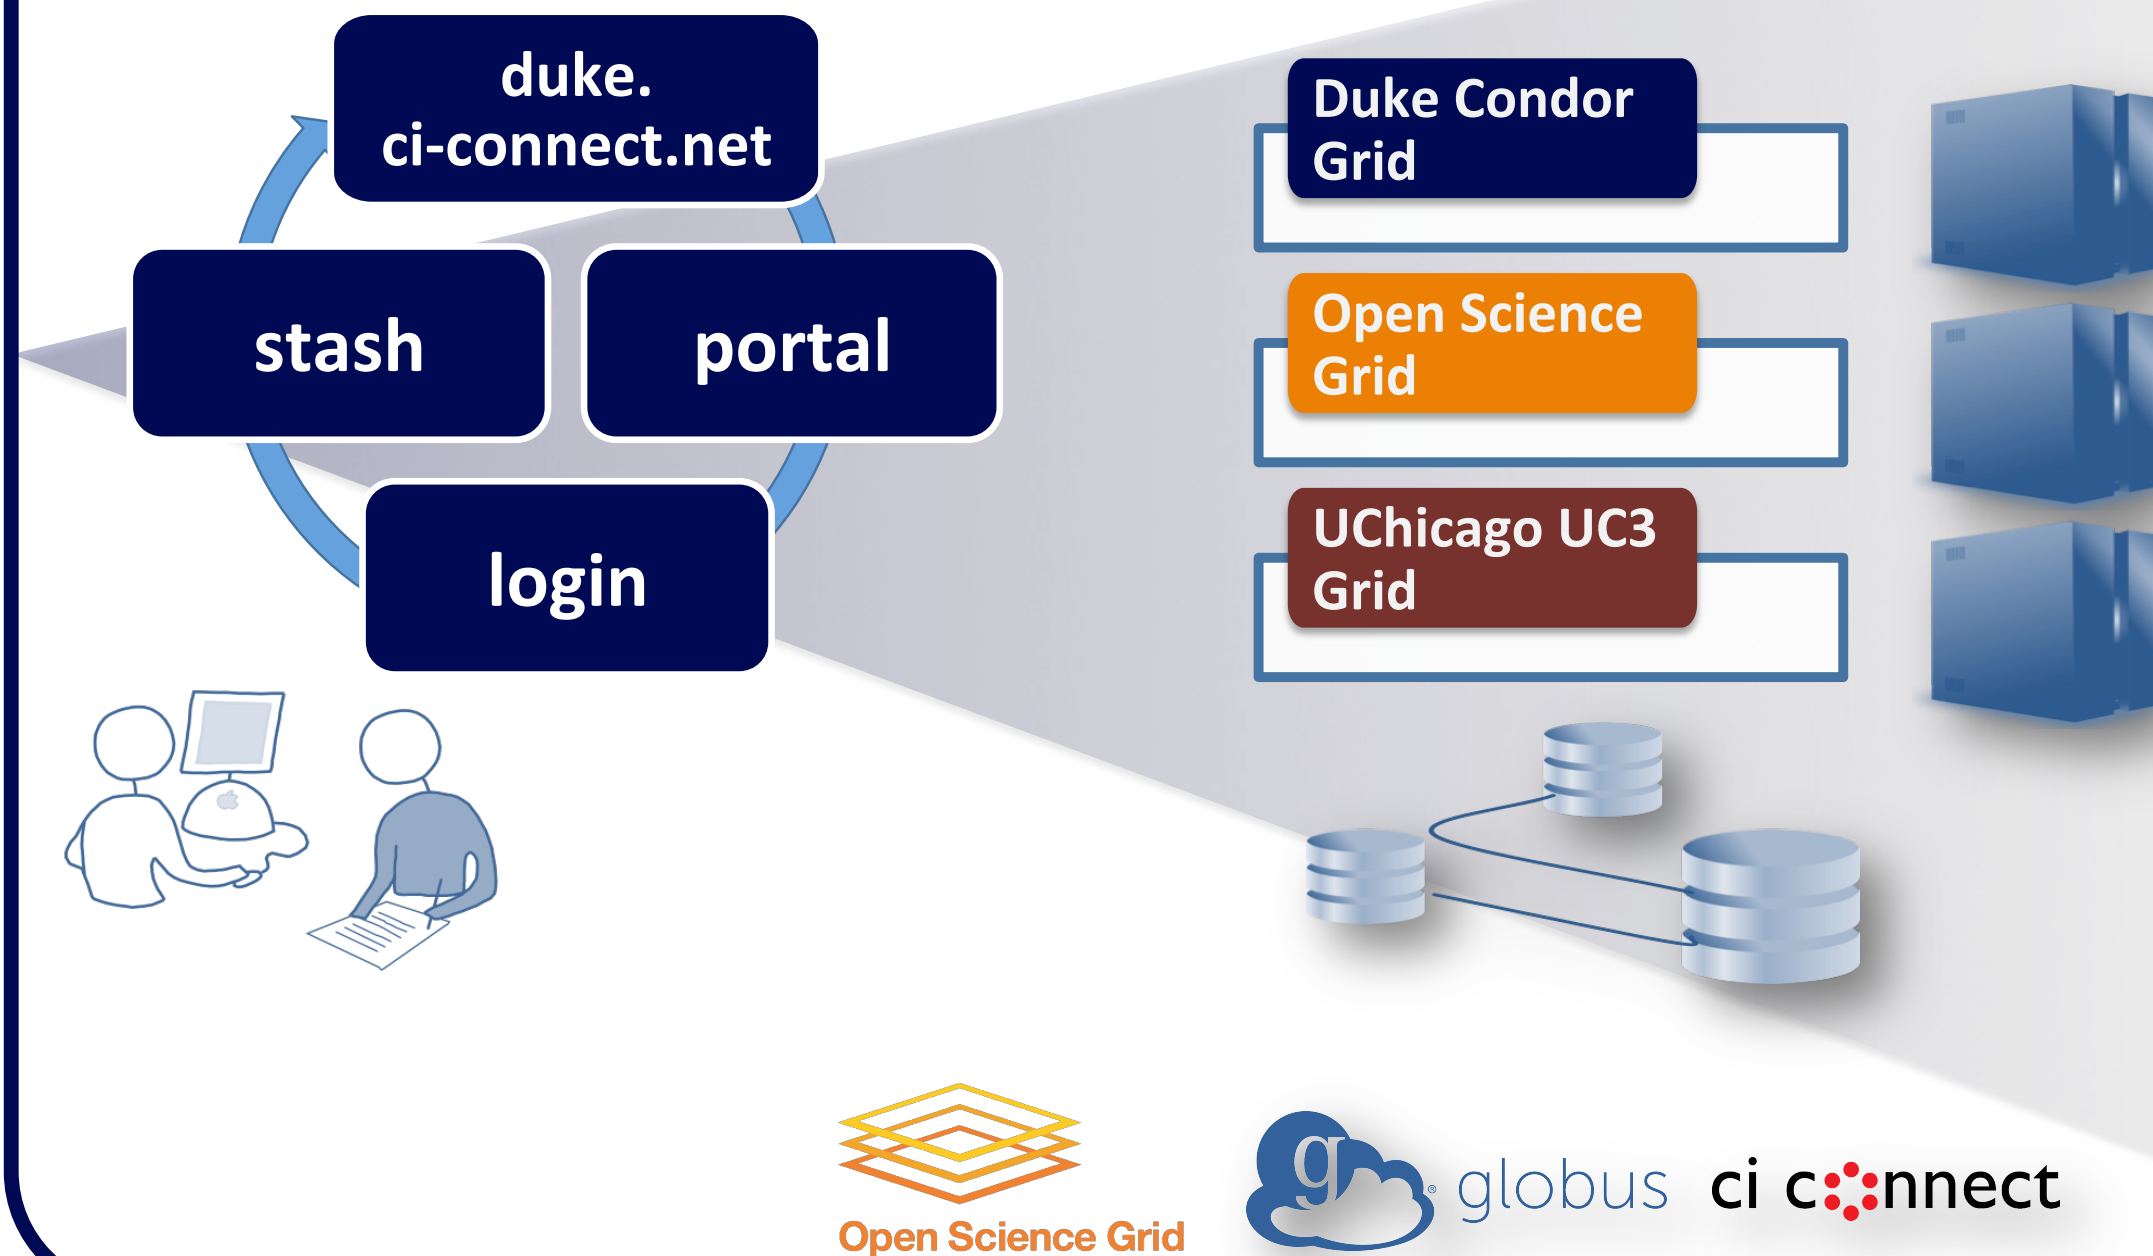
Components

- Leverages Globus, HTCondor, CI-Logon, U-Bolt, Bosco technologies
 - Bundled as instance of a CI Connect service portfolio
 - *Provided as a **Service** to reduce Campus IT load*
- Submit host
 - Flocks to OSG VO front-end, UC3 grid, & Amazon if needed
- Object storage service (90 TB usable)
 - POSIX, Globus Online, http, chirp access protocols
- Accounting (Gratia) and monitoring (Cycle Server) services



osg connect





Maturing Service - OASIS

- We've got about a full year of experience in running the current OASIS service.
 - A few operational hiccups, but has been getting a basic service to VOs.
- We're in the process of planning major improvements to this service.
 - Among other features, this will allow VOs to host external repositories. Users could do software installation from the "comfort of home" but publish easily to the OSG.

New Approach - Traceability

- One significant usability hurdle for new users has been acquiring and managing certificates and proxies.
- Getting a certificate, putting it in the browser, and transferring it to a login UI still is significant voodoo for new users.
- The security team re-evaluated the basic tenets of why we need certificates for users. This boiled down to one thing: *traceability*.



Traceability Project

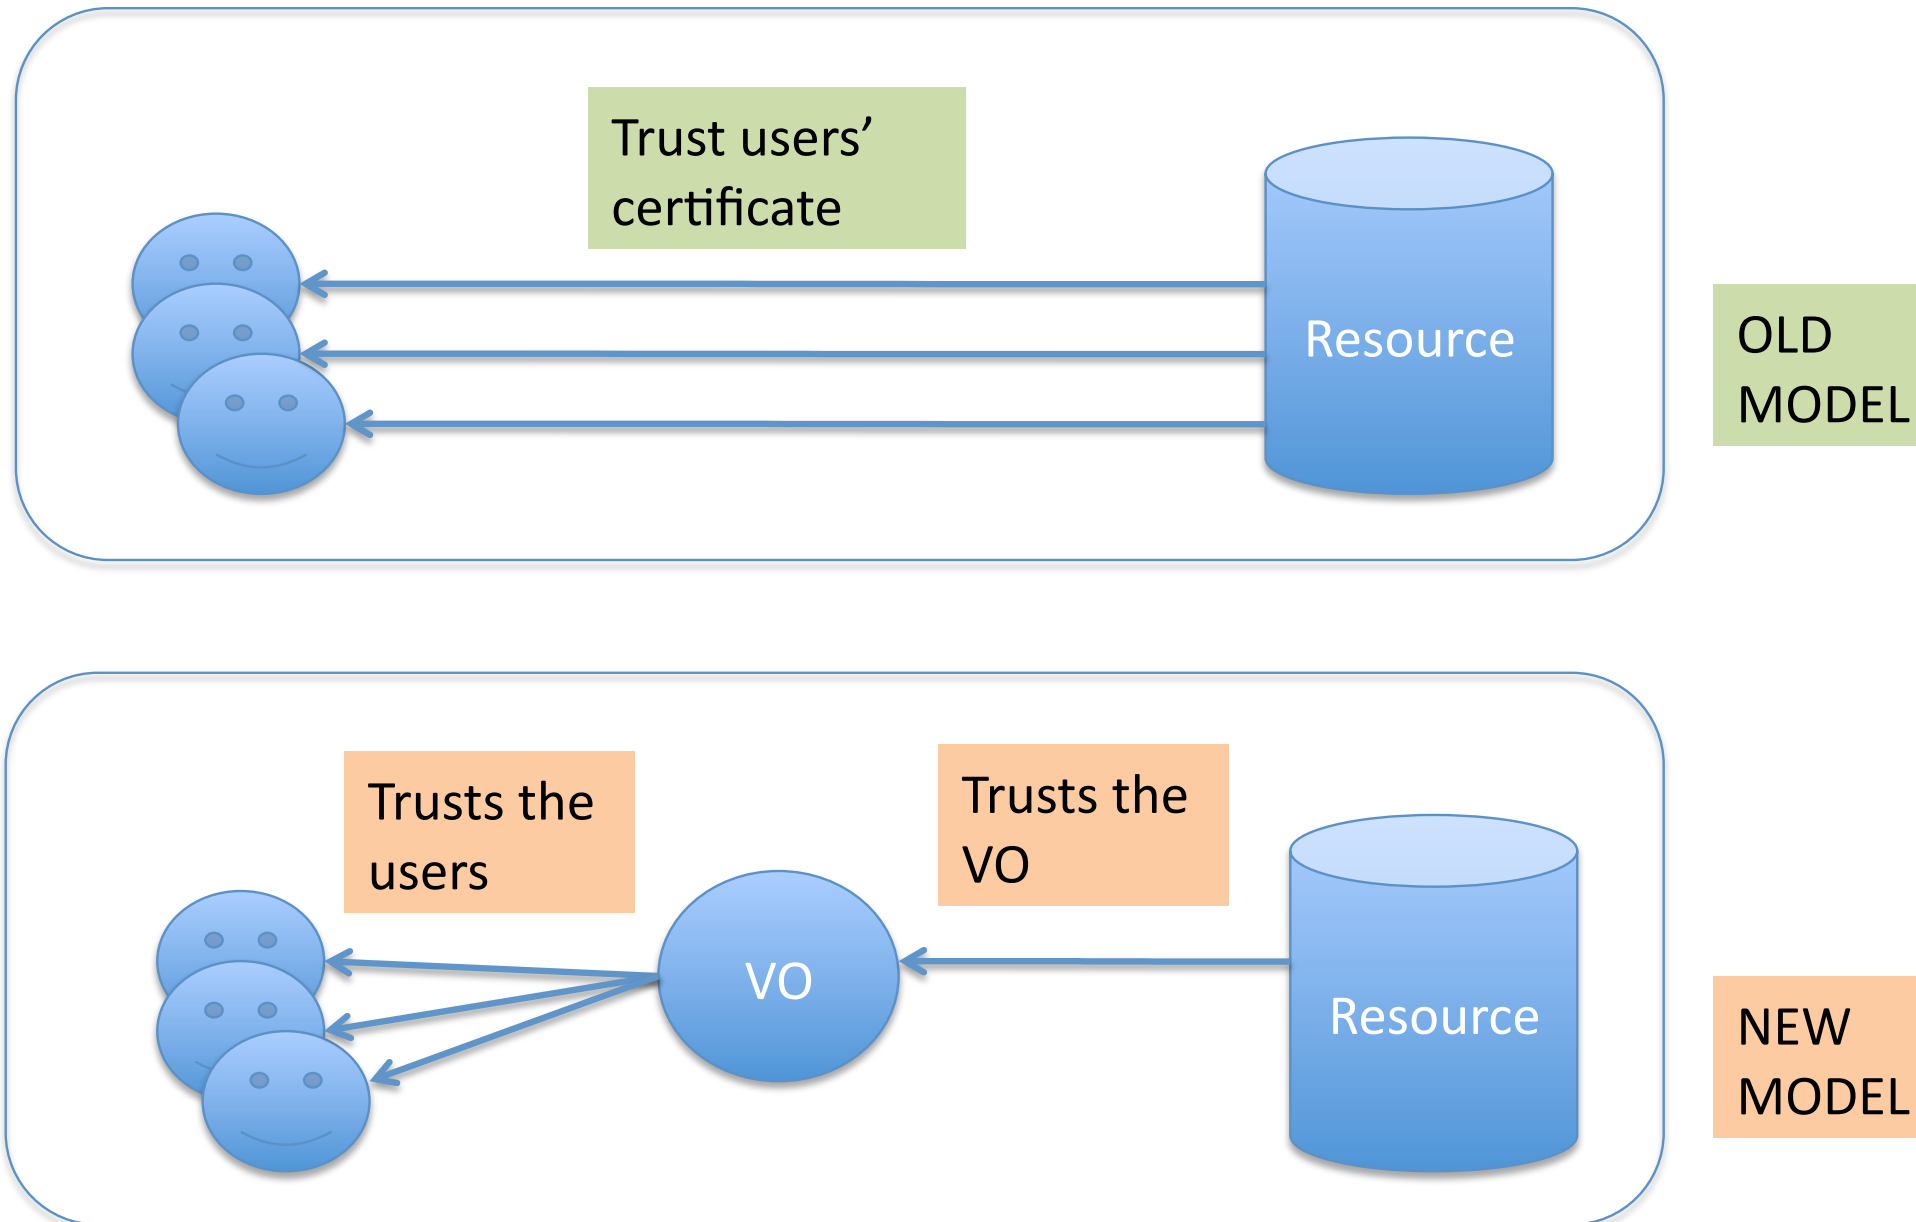
- Traceability of User Jobs: Goal is eliminating end user certificates
 - Traceability = associating users with their jobs
 - Who owns this job? Can we answer this question without certificates?
 - Proved that GlideinWMS system can trace user jobs even without certificates.
 - OSG-XSEDE VO and GLOW VO are the first beneficiaries. Evaluated their user management practices and job submission systems

Slide Courtesy of Mine Altunay



Open Science Grid

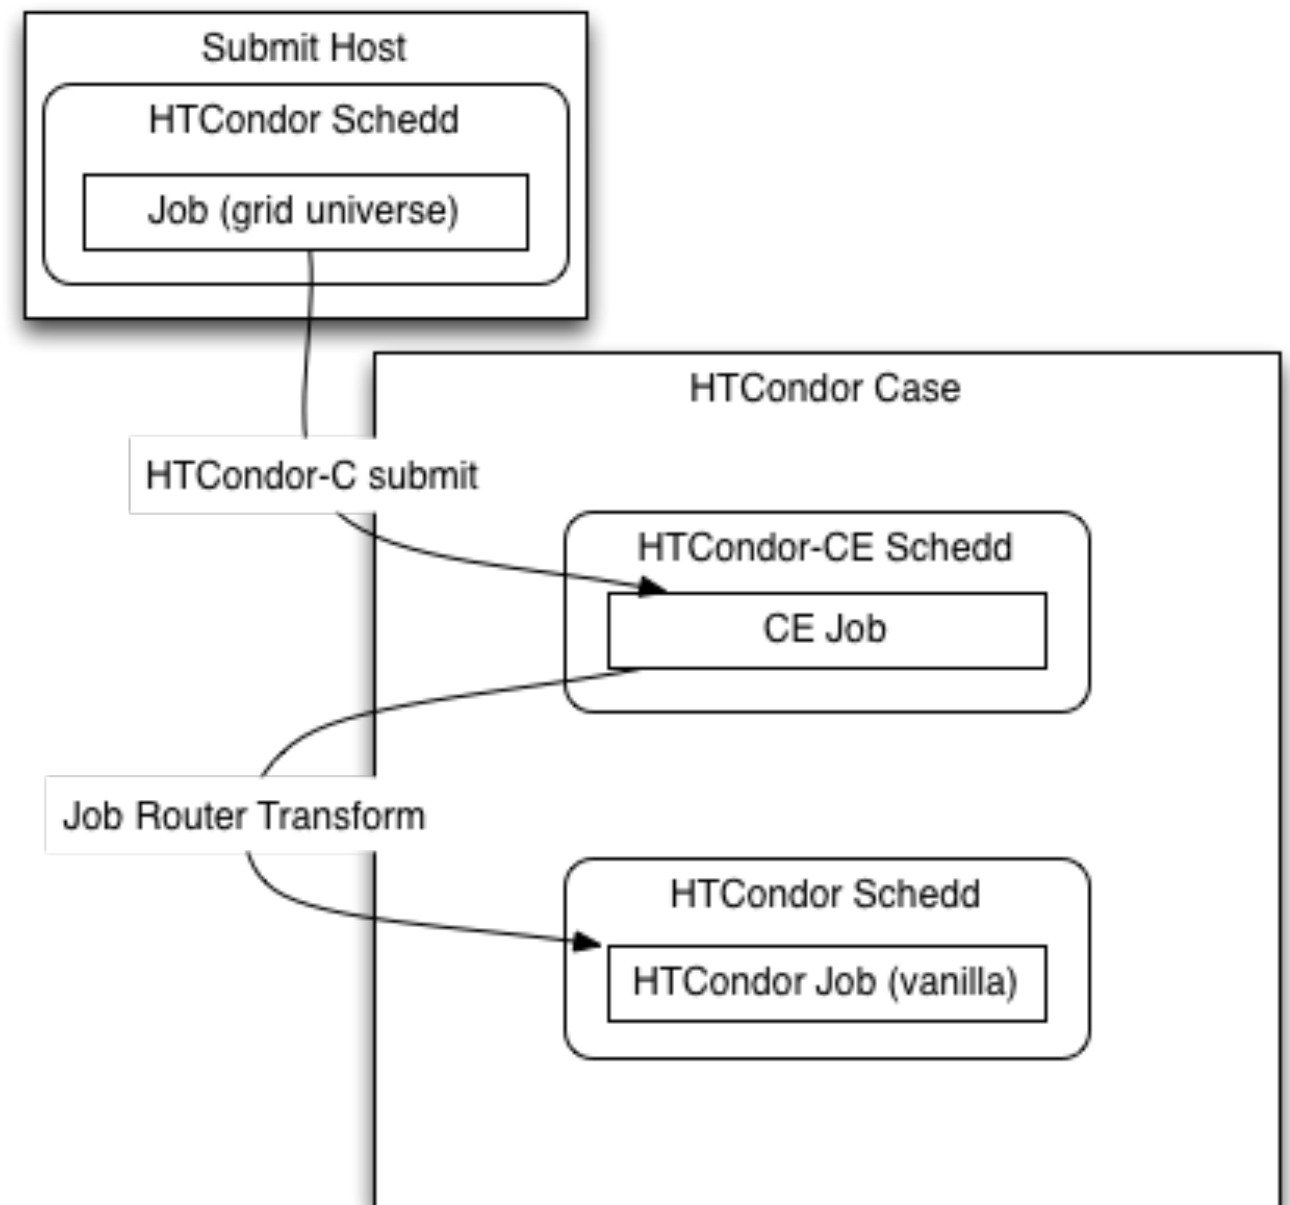
Traceability Project: Changing Trust Relationships



Slide Courtesy of Mine Altunay.

New Components - HTCondor-CE

- OSG 3.2 features the HTCondor-CE as OSG's next generation gatekeeper technology.
- HTCondor-CE should be more scalable, more robust, and (most importantly) easier to debug.



HTCondor-CE

```
bbockelm — root@red:~ — ssh — 92x24

360622.0  uscmsPool2557  4/9  07:23  0+00:13:02 R  0  0.0  glidein_startup.sh
360623.0  uscmsPool2553  4/9  07:23  0+00:00:00 I  0  0.0  glidein_startup.sh
360624.0  uscmsPool2556  4/9  07:29  0+00:00:00 I  0  0.0  glidein_startup.sh
360625.0  uscmsPool2558  4/9  07:30  0+00:00:00 I  0  0.0  glidein_startup.sh
360626.0  uscmsPool2558  4/9  07:30  0+00:00:00 I  0  0.0  glidein_startup.sh
360627.0  uscmsPool2555  4/9  07:30  0+00:00:00 I  0  0.0  glidein_startup.sh
360628.0  uscmsPool2557  4/9  07:34  0+00:00:00 I  0  0.0  glidein_startup.sh
360629.0  uscmsPool2557  4/9  07:34  0+00:00:00 I  0  0.0  glidein_startup.sh
360630.0  uscmsPool2557  4/9  07:34  0+00:00:00 I  0  0.0  glidein_startup.sh
360631.0  uscmsPool2557  4/9  07:34  0+00:00:00 I  0  0.0  glidein_startup.sh
360632.0  uscmsPool2557  4/9  07:34  0+00:00:00 I  0  0.0  glidein_startup.sh
360633.0  uscmsPool2557  4/9  07:37  0+00:00:00 I  0  0.0  glidein_startup.sh
360634.0  uscmsPool2557  4/9  07:37  0+00:00:00 I  0  0.0  glidein_startup.sh
360635.0  uscmsPool2557  4/9  07:37  0+00:00:00 I  0  0.0  glidein_startup.sh
360636.0  uscmsPool2557  4/9  07:37  0+00:00:00 I  0  0.0  glidein_startup.sh
360637.0  uscmsPool2557  4/9  07:41  0+00:00:00 H  0  0.0  glidein_startup.sh
360638.0  uscmsPool2557  4/9  07:41  0+00:00:00 H  0  0.0  glidein_startup.sh
360639.0  uscmsPool2557  4/9  07:41  0+00:00:00 H  0  0.0  glidein_startup.sh
360640.0  uscmsPool2557  4/9  07:41  0+00:00:00 H  0  0.0  glidein_startup.sh
360641.0  uscmsPool2557  4/9  07:41  0+00:00:00 H  0  0.0  glidein_startup.sh
360642.0  uscmsPool2557  4/9  07:41  0+00:00:00 H  0  0.0  glidein_startup.sh

3090 jobs; 10 completed, 2 removed, 1110 idle, 1481 running, 487 held, 0 suspended
[root@red ~]# condor_ce_q | tail -n 24
```

<https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallHTCondorCE>

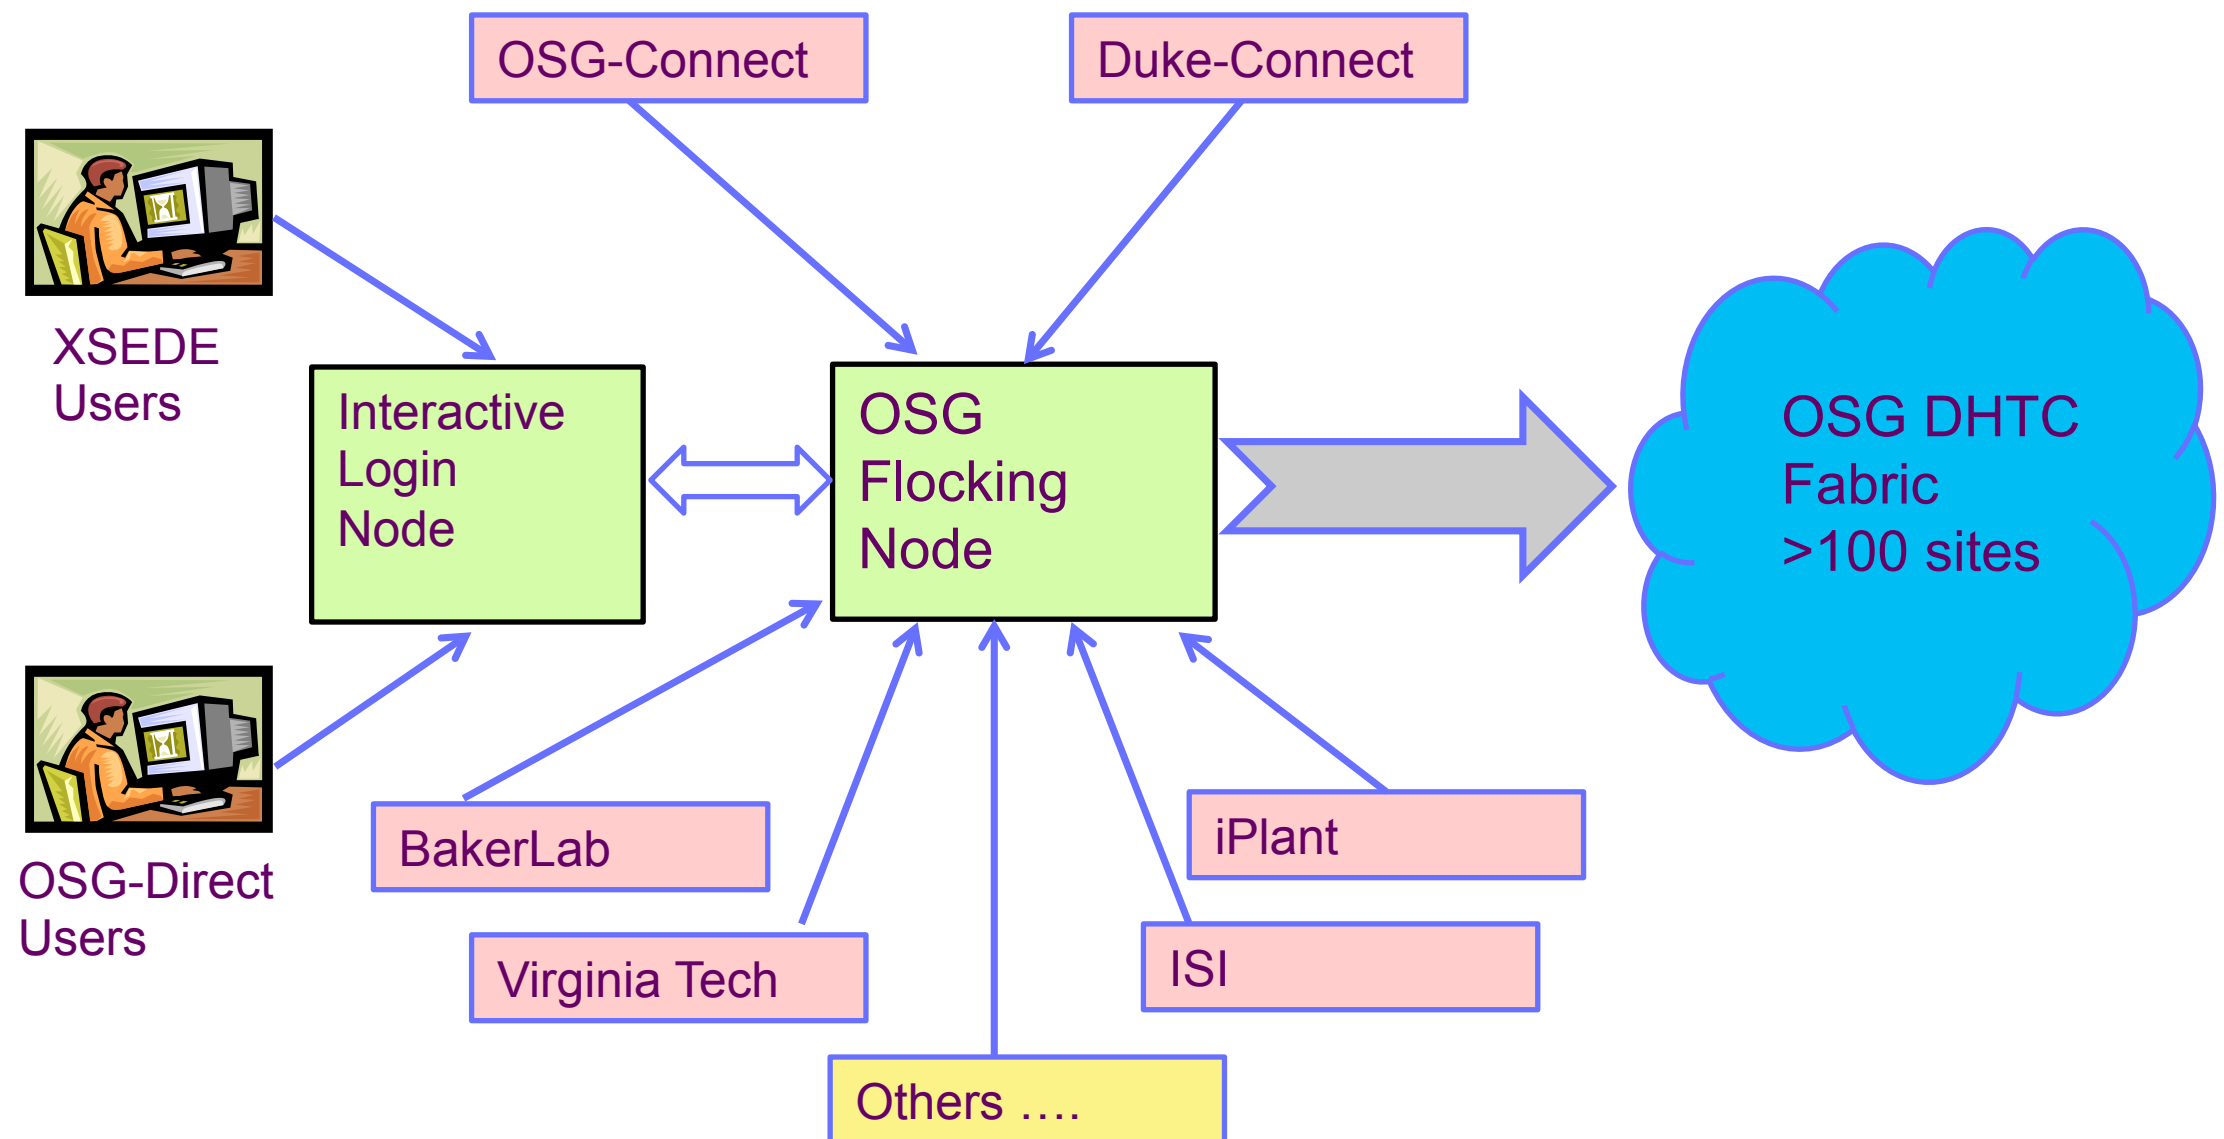
New Components - HTCondor-CE

- The first preview of HTCondor-CE was released almost 12 months ago.
- Ramp-up has been slow, largely because we had to wait for client components to add support.
- As of March, we have a fairly robust release that anyone should be able to use. I recommend this as the default for anyone who is updating their CE.

New (Old) Use Cases

- One of the big projects for the next year is to reinvent the osg-client.
- The current OSG client (and a majority of documentation) is from the pre-pilot era.
- We would like to package a submit node install for sites who would like to connect to the OSG VO.
 - Right now, flocking to the OSG VO is a process - a long checklist - not a product you can install.
- Otherwise, individual users will be steered to OSG Connect.

Access to OSG DHTC Fabric via OSG VO

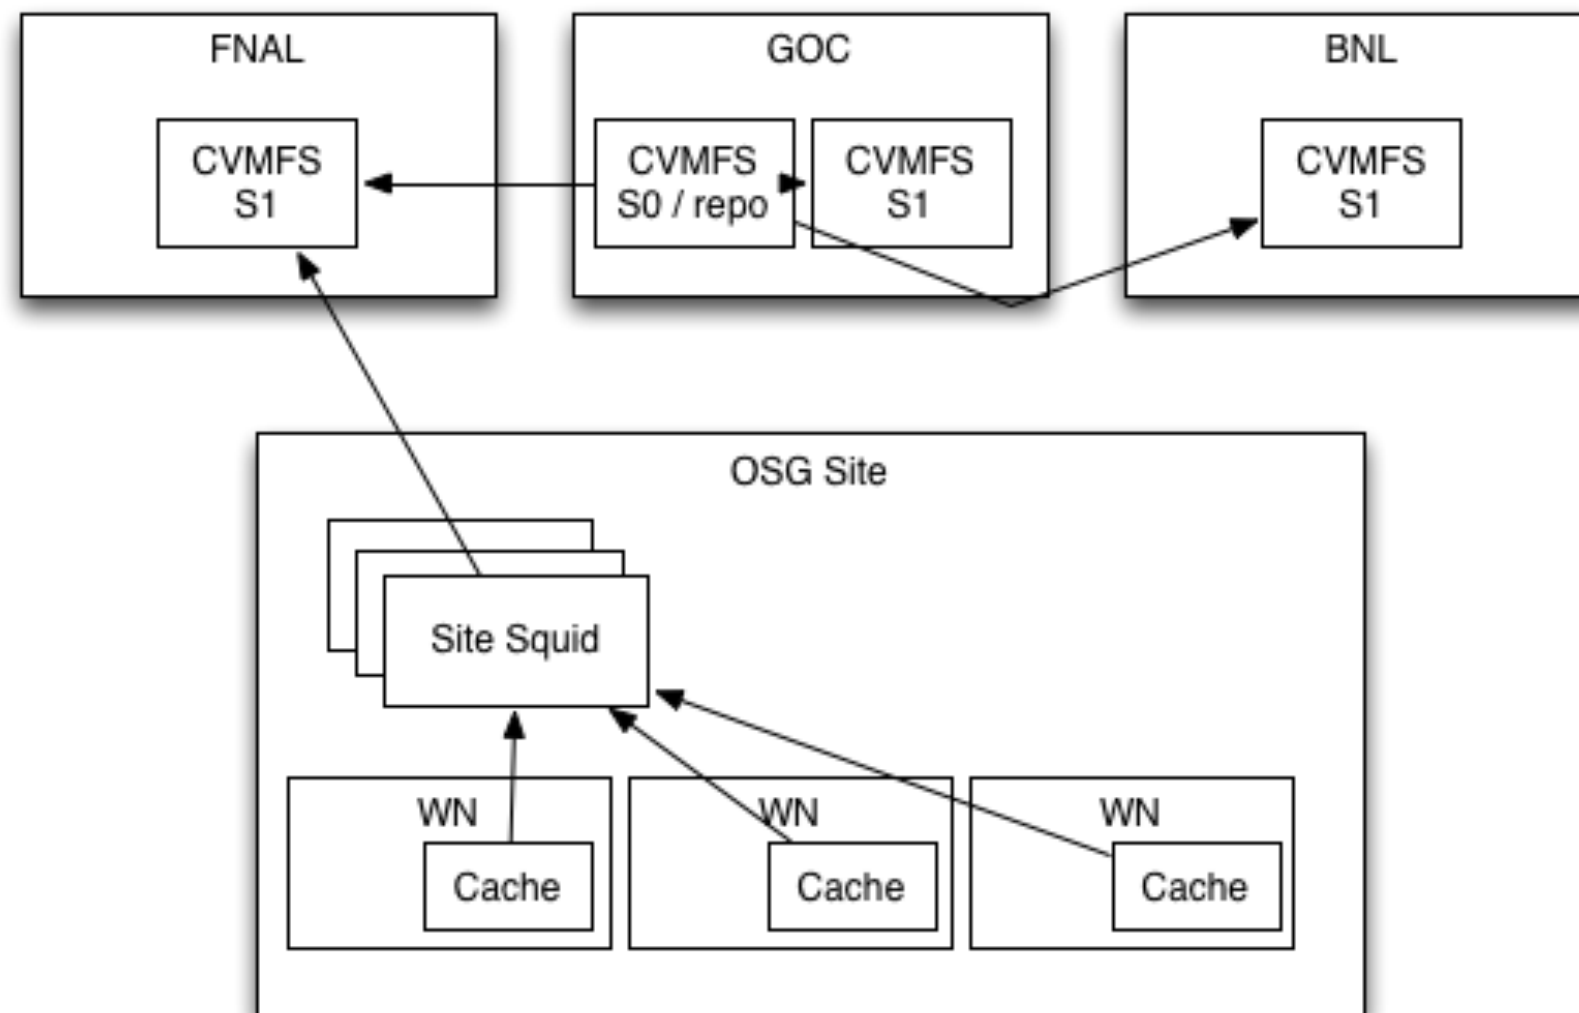


All access operates under the OSG VO using glideinWMS

The Data Question

- We have built second or third generation products on top of HTCondor to help users run jobs on the OSG Production Grid. What about data?
- This year, we pushed Squid / CVMFS to its current set of limits.
 - CVMFS does a fantastic job in helping users create a portable application, especially when combined with Parrot for non-CVMFS sites.
 - It is very sensitive to the *working set size* - the volume of data each job will touch and the volume of data several jobs will touch. It does well at software distribution - where the working set size is often <500MB, but poorly at data distribution - where the working set size is >1GB.
- I think the Next Big Thing OSG will try to tackle is the case where every job in a workflow needs the same 10GB of the input.

What Exists - OASIS/CVMFS



OASIS works well for *software* distribution, but not currently for data. Limitations are mostly due to the Squid size and cache size.

Where Next?

- This isn't clear!
- Options include:
 - Working with sites to expand the CVMFS infrastructure.
 - Using “alien caches” to keep the CVMFS cache on a larger shared file system.
 - Wider rollout of a different technology - iRODS / OSG Public Storage.
 - Else?

Conclusions

- Not all problems are technological; what I discussed today are only a small portion of the OSG Fabric of Services.
- The Software and Release teams provide support for software in various states of their lifecycle - from first production release to mature to deprecated to orphaned.
 - The existing “Software Factory” keeps this set of software as a coherent, well-tested distribution.
- OSG Connect is a new initiative for an old problem - how does one best bootstrap DHTC at a campus?
- OSG Security’s traceability project allows interested VOs to decrease the need. As we go forward, we will eliminate more use cases for long-term certificates.
- “Simple” data management remains not-so-simple and will be a top priority in the next year.